

APPLICATION  
FOR  
UNITED STATES LETTERS PATENT

TITLE: DISTRIBUTED IMAGE STORAGE ARCHITECTURE  
APPLICANT: GARY TESSMAN, JR. and PATRICK LIPPERT

10007695-121001

## DISTRIBUTED IMAGE STORAGE ARCHITECTURE

This application claims priority to U.S. provisional application 60/251,834, filed December 8, 2000, which is incorporated by reference in its entirety.

5

### TECHNICAL FIELD

The present invention relates generally to an online communications system and more particularly to an online communications system for processing digital images.

10

### BACKGROUND

In the client-server network architecture of the Internet and/or Web, electronic documents are stored in computer systems running server programs and are accessed by computer systems running client programs. For example, information on the Web is provided by Web servers and is accessible by a client program such as a Web browser (e.g., Netscape's Navigator, Microsoft's Internet Explorer, Java's micro-browser).

15

Information on the Internet and/or Web may be represented by specially formatted text files (e.g., Web pages) written in Hypertext Markup Language ("HTML") or some other markup language, such as XML, HDML, and/or VRML. Each text file may be identified by a network address such as a Universal Resource Locator ("URL"). A typical Web page may include one or more hyperlinks referring to the network addresses of other Web pages. Hyperlinks may be displayed as underlined text or graphical images that, when clicked, send a request for the associated page. For example, when a hyperlink in a home page is selected, a request is sent out by the client to the address specified in the hyperlink, and the associated Web page is downloaded and displayed, replacing the home page on the user's screen with the associated Web page.

20

25

Browsers and other client programs typically use a communications protocol such as Hypertext Transfer Protocol ("HTTP") to request pages from Web servers. HTTP is a request/response protocol. Through a connection established between a client and a server, the client sends a request to the server, and the server provides a response to the client.

30

An Internet service provider ("ISP") may be used to provide subscribers with access to the Internet and/or World Wide Web ("Web"). In general, a subscriber relies on an ISP to

provide computers that are connected to and therefore enable communication over the Internet and/or Web. An ISP may offer services in addition to basic Internet access such as, for example, providing e-mail and instant messaging services enabling electronic communication, Web-hosting services allowing subscribers to publish homepages, newsgroup services

5 allowing subscribers to read and post to newsgroups, and image services allowing subscribers to view and order digital images of pictures from a developed film roll.

### SUMMARY

In one general aspect, a computer system stores digital images within a computer

10 system by identifying a first storage facility and a directory within the first storage facility for storing a digital image; generating a first image identifier associated with the first storage facility and the directory; generating a second image identifier comprising a random number; generating a unique hash value by encrypting the first and second image identifiers; and identifying a storage path using the first and second image identifiers and the unique hash

15 value such that related digital images have unrelated storage paths. The unique hash value may be generated by applying the MD5 algorithm or the DEC algorithm to the first and second image identifiers.

Implementations may include one or more of the following features. For example, the digital image may be received from at least one of a subscriber of the computer system and a

20 third party associated with the subscriber. Identifying a storage path may include extracting and/or translating storage path information from at least one of the first image identifier and the second image identifier. Identifying a storage path also may include using the unique hash value as a filename. The digital image may be stored in the first storage facility at the identified storage path along with Lower resolution thumbnails.

Other general aspects may include identifying a second storage facility for storing

25 metadata describing the digital image. Identifying the second storage facility may include encoding account information, such as a screen name, associated with the digital image. The encoded account information may be mapped to an appropriate storage space group containing the second storage facility. Access may be provided to the stored digital image and

30 the stored metadata.

Aspects of the present invention may be implemented by an apparatus and/or by a computer program stored on a computer readable medium. The computer readable medium may comprise a disc, a client device, a host device, and/or a propagated signal.

Other features and advantages will be apparent from the following description,  
5 including the drawings, and from the claims.

### DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram illustrating aspects of a computer system.

Fig. 2 is a block diagram expansion of aspects of Fig. 1.

10 Fig. 3 is a block diagram expansion of aspects of Fig. 2.

Fig. 4 is a block diagram expansion of aspects of Fig. 3.

Fig. 5 is a flowchart of a method that may be implemented by the computer system illustrated in Figs. 1-4.

15 Fig. 6 is a flowchart of another method that may be implemented by the computer system illustrated in Figs. 1-4.

### DETAILED DESCRIPTION

Fig. 1 illustrates an exemplary computer system 100 for implementing techniques to process digital images. For brevity, several elements in the figure are represented as  
20 monolithic entities. However, these elements each may include numerous interconnected computers and components designed to perform a set of specified operations.

As shown, the computer system 100 includes a client system 110 connected through a network 15 to a host system 20. The client system 10 is configured to send requests and the host system 20 is configured to respond to requests. The host system 20 may include and/or  
25 form part of an information delivery network, such as, for example the Internet, the World Wide Web, an online service provider, and/or any other analog or digital wired and/or wireless network that provides information. Such an information delivery network may support a variety of online services including Internet and/or web access, e-mail, instant messaging, paging, chat, interest group, audio and/or video streaming, and/or directory services.

30 In general, the client system 10 includes a computer system having hardware and/or software components for communicating with the network 15 and the host system 20. The

client system 10 and host system 20 each may include one or more general-purpose computers (e.g., personal computers and/or servers), one or more special-purpose computers (e.g., devices specifically programmed to communicate with each other), or a combination of one or more general-purpose computers and one or more special-purpose computers. The client

5 system 10 and host system 20 may be structured and arranged to communicate using various communication protocols (e.g., http, WAP) and encapsulation protocols (e.g., UDP) to establish connections (e.g., peer-to-peer) between network elements and/or to operate within or in concert with one or more other systems (e.g., the Internet and/or Web).

In one implementation, the client system 10 and the host system 20 each include a

10 device (e.g., client device 12, host device 22) operating under the command of a controller (e.g., client controller 14, host controller 24). An example of a device is a general-purpose computer capable of responding to and executing instructions in a defined manner. Other examples include a special-purpose computer, a personal computer ("PC"), a workstation, a server, a laptop, a Web-enabled telephone, a Web-enabled personal digital assistant ("PDA"),

15 an interactive television set, a set top box ("STB"), video tap recorder ("VTR"), a digital video disc ("DVD") player, or any other component, machine, tool, equipment, or some combination thereof capable of responding to and executing instructions.

An example of a controller is a software application (e.g., operating system, browser application, microbrowser application, server application, proxy application, gateway

20 application, tunneling application, e-mail application, instant messaging client application, online service provider client application, interactive television client application, and/or ISP client application) loaded on a device commanding and directing communications enabled by the device. Other examples include a computer program, a piece of code, an instruction, another device, or some combination thereof, for independently or collectively instructing the

25 device to interact and operate as desired. The controller may be embodied permanently or temporarily in any type of machine, component, physical or virtual equipment, storage medium, or propagated signal capable of providing instructions to a device. In particular, the controller (e.g., software application and/or computer program) may be stored on a storage media or device (e.g., read only memory ("ROM"), magnetic diskette, or propagated signal)

30 readable by a general or special purpose programmable computer, such that the functions described herein are performed if the storage media or device is read by a computer system.

The network 15 may include one or more delivery systems for directly or indirectly connecting the client system 10 and the host system 20, irrespective of physical separation. Examples of delivery systems include, but are not limited to, a local area network ("LAN"), a wide area network ("WAN"), the Internet, the Web, a telephony network (e.g., analog, digital, 5 wired, wireless, PSTN, ISDN, or xDSL), a radio network, a television network, a cable network, a satellite network, and/or any other wired or wireless communications network configured to carry data. Each network may include one or more elements, such as, for example, intermediate nodes, proxy servers, routers, switches, adapters, and wired or wireless data pathways, configured to direct and/or deliver data.

10 In one implementation, the client system 10 includes a computer system running an ISP client application. The client system 10 allows a subscriber to access online content and receive a variety of online services (e.g., Internet access, e-mail, chat, newsgroup access and/or instant messaging). In particular, the client system 10 allows a subscriber to view digital images from a processed (e.g., developed) roll of film. The host system 20 enables a 15 subscriber to download, access, send, share, and receive digital images from a processed roll of film.

Referring to Fig. 2, the communications system 200 is an expansion of the block diagram of Fig. 1, focusing primarily on one particular implementation of the host system 20. The host system 20 includes a host device 22 and a host controller 24. The host controller 24 20 generally is capable of transmitting instructions to any or all of the elements of the host device 22. For example, in one implementation, the host controller 23 includes one or more software applications loaded on the host device 22. In other implementations, the host controller 24 may include any of several other programs, machines, and devices operating independently or collectively to control the host device 22.

25 The host device 22 includes a login server 210 for enabling access by subscribers and for routing communications between the client system 10 and other elements of the host device 22. The host device 22 also may include various host complexes, such as the depicted Instant Messaging ("IM") host complex 220 and Online Service Provider ("OSP") host complex 230. To enable access to these host complexes by subscribers, the client system 10 30 includes communication software, for example, an IM client application, an OSP client application, and/or a browser application. The IM and OSP communication software

applications are designed to facilitate the subscriber's interactions with the respective services and, in particular, may provide access to all the services available within the respective host complexes. The login server 210 may implement one or more authorization procedures to enable simultaneous access to the IM host complex 220 and the OSP host complex 230. The

5 IM host complex 220 and the OSP host complex 230 may be connected through one or more gateways (not shown) that perform protocol conversions necessary to enable communications between the IM host complex 220, the OSP host complex 230, and the Internet 30.

The IM host complex 220 generally is independent of the OSP host complex 230, enabling support of instant messaging services irrespective of a subscriber's network or

10 Internet access. Thus, the IM host complex 220 allows subscribers to send and receive instant messages, whether or not they have access to any particular ISP. The IM host complex 220 may support associated services, such as administrative matters, advertising, directory services, chat, and interest groups related to instant messaging. The IM host complex 220 has an architecture that enables all of the machines within the IM host complex to communicate

15 with each other. To transfer data, the IM host complex 220 employs one or more standard or exclusive IM protocols.

Typically, the OSP host complex 230 supports different services, such as e-mail services, discussion group services, chat, news services, and Internet access. In addition, the

20 OSP host complex 230 may offer instant messaging services to its subscribers independent of or based on IM host complex 220. The OSP host complex 230 generally is designed with an architecture that enables the machines within the OSP host complex 230 to communicate with each other and employs certain protocols (i.e., standards, formats, conventions, rules, and structures) to transfer data. For instance, the OSP host complex 230 ordinarily employs one or more OSP protocols and custom dialing engines to enable access by selected client

25 applications. The OSP host complex 230 may define one or more specific protocols for each service based on a common, underlying proprietary protocol.

The OSP host complex 230 supports a set of services from one or more servers located internal to and external from the OSP host complex 230. For purposes of this discussion, servers external to the OSP host complex 230 generally may be viewed as existing on the

30 Internet 30. Servers internal to the OSP complex 230 may be arranged in one or more

configurations. For example, servers may be arranged in centralized or localized clusters in order to distribute servers and subscribers within the OSP host complex 230.

5 In the implementation shown by Fig. 2, the OSP host complex 230 includes a routing processor 232. In general, the routing processor 232 will examine an address field of a data request, use a mapping table to determine the appropriate destination for the data request, and direct the data request to the appropriate destination. More specifically, in a packet-based implementation, the client system 10 may generate information requests, convert the requests into data packets, sequence the data packets, perform error checking and other packet-switching techniques, and transmit the data packets to the routing processor 232. Upon 10 receiving data packets from the client system 10, the routing processor 232 may directly or indirectly route the data packets to a specified destination within or outside of the OSP host complex 230. For example, in the event that a data request from the client system 10 can be satisfied locally, the routing processor 230 may direct the data request to a local server 234. In the event that the data request cannot be satisfied locally, the routing processor 232 may direct 15 the data request externally to the Internet 30 or the IM host complex 220.

20 The OSP host complex 230 also includes a proxy server 236 for directing data requests and/or otherwise facilitating communication between the client system 10 and the Internet 30. The proxy server 236 may include an Internet Protocol ("IP") tunnel for converting data between an OSP protocol and standard Internet protocol to enable the client system 10 to communicate with the public Internet 30.

25 The proxy server 236 also may allow the client system 10 to use standard Internet protocols and formatting to access the OSP host complex 230 and the Internet 30. For example, the subscriber may use an OSP TV client application having an embedded browser application installed on the client system 10 to generate a request in standard Internet protocol, such as HyperText Transport Protocol ("HTTP"). In a packet-based implementation, data packets may be encapsulated inside a standard Internet tunneling protocol, such as, for example, User Datagram Protocol ("UDP") and routed to the proxy server 236. The proxy server 236 may include an Layer Two Tunneling Protocol ("L2TP") tunnel capable of establishing a point-to-point protocol (PPP) session with the client system 10.

30 The proxy server 236 also may act as a buffer between the client system 10 and the Internet 30, and may implement content filtering and time saving techniques. For example,

the proxy server 236 can check parental controls settings of the client system 10 and request and transmit content from the Internet 30 according to the parental control settings. In addition, the proxy server 236 may include one or more caches for storing frequently accessed information, or may enable access to similar caches stored elsewhere. If requested data is  
5 determined to be stored in the caches, the proxy server 236 may send the information to the client system 10 from the caches and avoid the need to access the Internet 30.

The OSP host complex 230 further includes a mail system 238, an image farm 240, a film handler 242, an account manager 244, and a monitoring system 246 in communication with each other as well as the other elements in the communications system 200.

10 The mail system 238 is configured to receive, store, retrieve, route, and deliver electronic mail ("e-mail") messages. In general, the mail system 238 includes a system of folders or mailboxes associated with the subscribers of the OSP host complex 230 and a massive storage area for storing the contents of e-mail messages including attachments to the e-mail messages. When the mail system 238 receives an e-mail message addressed to a  
15 particular subscriber, the mail system stores the content and attachments of the e-mail message, inserts a link (e.g., href) or pointer corresponding to the storage location into the subscriber's mailbox, and alerts the subscriber of the new mail. The subscriber opens the e-mail message by logging in to the mailbox and selecting an icon including the link to the stored e-mail message. Similarly, a subscriber may send an e-mail message by logging in to  
20 the mailbox, generating an e-mail message, and then selecting a "send" button that causes the mail system 238 to store and forward the e-mail message to one or more intended recipients.

The image farm 240 is configured to receive, store, retrieve, route, and/or deliver digital images from developed rolls of film. The image farm 240 includes a system of storage servers and databases for storing digital images and metadata describing the digital images.

25 Typically, digital images and metadata are stored in a distributed fashion across several storage facilities and/or databases. For example, digital images may be stored independently of the submitting subscriber and of the original film roll such that image data for a particular subscriber is maintained over several different storage facilities. Thus, the overall integrity of the image service provided by the OSP host complex 230 can be preserved even in the event  
30 of a catastrophic outage of several storage facilities. In addition, metadata describing the image data may be distributed across a system of databases such that no one database is

responsible for a disproportionate amount of metadata. The image farm 240 is described in greater detail below.

The film handler 242 is configured to receive and route digital images to the image farm 240 as well as update metadata maintained in the image farm 240. In general, upon receiving a collection of digital images, the film handler 242 transfers the image data to the image farm 240 for storage and updates metadata associated with the digital images. Typically, the film handler 242 receives digital images from a subscriber or from a third party (e.g., film developer) associated with the subscriber of the OSP host complex 230. For example, the film handler 242 may receive one or a collection of digital images created with a digital camera from a subscriber. The subscriber may upload the digital images using one or more software applications provided by the OSP host complex 230 or the digital camera manufacturer. Alternatively, the film handler 242 may receive a collection of digital images from a film developer that received an undeveloped roll of film from a subscriber, developed the roll of film, and created digital images from the developed roll.

In one implementation the film handler 242 includes application server logic (e.g., Netscape Application Server (NAS) logic or Java Server (KJS) logic) for maintaining a pool of database connections for high performance accesses into databases in the image farm 240. The film handler 242 also may include a hash value-to-database function library for determining where database reads and writes occur within the databases. While the film handler 242 is depicted in Fig. 2 as being external to the image farm 240, in another implementations, the film handler 242 may reside as an application (e.g., Java applet) running within the image farm 240.

The account manager 244 is configured to maintain an image service account for the subscribers of the OSP host complex 230. In general, the account manager 244 creates new accounts by requesting information from and inserting submitted account information into account tables maintained in the image farm 240. In one implementation the account manager 244 includes application server logic (e.g., Netscape Application Server (NAS) logic or Java Server (KJS) logic) for maintaining a pool of database connections for high performance accesses into databases in the image farm 240. The account manager 244 also may include a hash value-to-database function library for determining where database reads and writes occur within the databases. In general, the account manager 244 creates an image service account

and populates a certain account table in the image farm 240. The account manager 244 also may automatically generate certain account information (i.e., fields) based on the information submitted from subscribers. In addition, the account manager 244 may be configured to authenticate subscribers logging in to the image service, deliver path information (e.g., linking information) to subscribers in response to requests for various image service features, and transmit purging information in response to complaints received from subscribers.

The monitoring system 246 is configured to locate and remove stored images that violate policies established by administrators of the OSP host complex 230. In general, the policies regulate the type of content that is acceptable for sharing among subscribers. While the OSP host 230 provides subscribers with the ability to share digital images with other subscribers, the monitoring system 246 is designed to temper that ability in order to protect subscribers from viewing offensive content. Although the monitoring system 246 may act proactively by filtering submitted digital images and deleting digital images that are deemed offensive, the monitoring system 246 typically is configured to receive notifications of offending content from subscribers and take appropriate actions in response. The monitoring system may take different levels of action depending upon the nature of the offending content. Such actions may range, for example, from blocking the complaining subscriber from reviewing the offending digital image to deleting the digital image, canceling the offender's account, and alerting law enforcement authorities. The monitoring system 246 is described in greater detail below.

Fig. 3 is one implementation of the image farm 240 of Fig. 2. As shown, the image farm 240 includes a distribution server 2402 in communication with several image farm databases 2404. Each image farm database 2404 includes several tables (e.g., film and image tables 2406 and account tables 2408) for maintaining information associated with received digital images and the subscribers of the OSP host complex 230. In general, the distribution server 2402 is configured to identify storage locations within the image farm databases 2404. The distribution server 2402 is in communication with and capable of receiving information (e.g., image metadata) from various other elements of the OSP host complex 230 of Fig. 2 including the mail system 238, the film handler 242, the account manager 244, and the monitoring system 246. The distribution server 2402 also is configured to insert, update, and/or delete information from identified storage locations. For example, the distribution

server 2402 may receive metadata from the film handler 242, identify a storage location in the image farm database 2404 based on the metadata, and insert the metadata into an appropriate table at the identified storage location. The distribution server 2402 also may receive metadata from the monitoring system 246, identify a storage location in the image farm database 2404 based on the metadata, and delete/replace stored metadata from an appropriate table at the identified storage location.

The image farm 240 further includes an image storage server system 2410 having an image storage server 2412, an image write server 2414, an image read server 2416, and an image purge server 2418. In general, the image write server 2414 is configured to insert image data into the image storage server 2412, the image read server 2416 is configured to retrieve image data from the image storage server 2412, and the image purge server 2418 is configured to delete image data from the image storage server 2412.

The image storage server system 2410 is in communication with and capable of responding to various elements of Fig. 2, including the client system 10, the film handler 242, the account manager 244, and the monitoring system 246. Servers (e.g., image write server 2414, image read server 2416, and image purge server 2618) of the image storage server system 2410 may function in response to a communication from the client system 10, the film handler 242, the account manager 244, and/or the monitoring system 246. For example, the image write server 2414 may generate metadata and store image data in the image storage server 2412 in response to a communication from the film handler 242. The image read server 2416 may retrieve image data in response to a communication from a client system 10 according to path information provided by the account manager 244. The image purge server 2418 may delete image data in response to a communication from the monitoring system 246.

In one implementation, each image storage server 2412 may store digital images according to the following directory structure:

host:/data/AAA/BB/CC/DD/EE/ID\_N.type

In this pneumonic directory structure, host identifies a particular image storage server 2412 (root directory) for storing a digital image, data identifies a static directory of the image storage server 2412, AAA is a three hexadecimal character directory identifier denoting a

main storage directory of the image storage server 2412, BB, CC, DD, and EE are each two hexadecimal character subdirectories in the image storage server 2412, ID is an encrypted 32 hexadecimal character image identifier that provides a file name and contains location information, N is the largest pixel dimension of an image (width or height), and type identifies image type (e.g., jpg).

By design, the above directory structure limits the number of files that can exist in one subdirectory while still establishing enough subdirectories so that the system will be relatively sparse. With this directory structure, for example, an image storage server 2412 will never have more than 256 images or 256 subdirectories (designated by two hexadecimal digits) within a single tier. The file systems are mounted below the server root directory at the /data/AAA level. This ensures that the AAA level as well as the host level will be relatively sparse, while the BB, CC, DD and EE levels will be heavily populated. In one implementation, IDs are varied over the BB and host range quickly. The ID field resides beneath the EE level, since ID is a unique identifier and only one file will have that unique path. The directory structure is very robust, enabling each file system to store in excess of two billion images, even if only IDs are considered. It should be noted that the specific numerical values provided are exemplary only.

Fig. 4 provides further detail regarding the image farm databases 2404 of Fig. 3. As shown, the image farm databases 2404 are separated into multiple storage space groups referred to as "buckets." In one implementation, the image farm databases 2404 are separated into sixteen relatively equal-sized buckets (bucket 0 – bucket 15). As described above, each image farm database 2404 maintains information associated with received digital images and the subscribers of the OSP host complex 230. This information is received by the image farm databases 2404 and is stored within the film and image tables 2406 and account tables 2408 in a distributed fashion. Namely, the information received by the image farm databases 2404 is relatively uniformly distributed across more than one or all buckets. Distributing information across multiple storage spaces areas in this way eliminates the need to scan across multiple databases for the most common image service operations offered by the OSP host complex 230.

In one implementation, relatively uniform distribution across all buckets is accomplished by applying an OSP proprietary hashing code to at least one field in the

received information. Application of the OSP proprietary hashing code returns a unique hash value. The hash value then is transformed into a bucket number corresponding to one of the relatively equal-sized buckets. The information is routed to the appropriate bucket, each bucket containing sets of film and image tables 2406 and account tables 2408 configured to

5 store the received information in appropriate fields.

Each of the bucketed image farm databases 2404 includes film and image tables 2406 and account tables 2408. In one implementation, the film and image tables 2406 include a film reference table 2406a containing sharing information related to a film, a film table 2406b for describing a collection of images, an image reference table 2406c containing information

10 showing that an image has been used to construct a film, an image table 2406d for describing a particular image, a film attributes table (not shown) for storing future or less common roll/album properties that are not included in the film table, an image attributes table (not shown) for storing future or less common properties that are not included in the image table, a film identification counter (not shown) that may be used to generate a unique film identifier

15 without the need for an external server, and an image copyright table (not shown) for flagging images that have copyright information bound to them. The account tables 2408 include an account information table 2408a for storing account information associated with a particular subscriber.

The film reference table 2406a may include information such as: film identifiers,

20 reference owner, shared by name, buddy name, film reference identifier, film type, number of images, title, share type, access control, modification date, first viewed date, and view flag. In one implementation, fields for the film identifiers, reference owner, shared by name, buddy name, and film reference identifier are designated as primary keys for sorting and searching the film reference table as well as for establishing relationships among the film reference table

25 and other tables and/or databases.

The film table 2406b may include information such as: film identifiers, owner name, film type, number of images, creation status, creation date, modification date, expiration date, external film identifier, owner key, reference counter, access control, purge flags, title, description, background pattern, layout, or other attributes. In one implementation, fields for

30 the film identifiers are designated as primary keys for sorting the film table as well as for establishing relationships among the film table and other tables and/or databases.

The image reference table 2406c may include information such as: film identifiers, image identifiers, image version, image counter, sequence number, image hash value, caption, creation date, modification date, frame type, image type, resolution, maximum width, maximum height, access control, purge flags, and branding identifier. In one implementation, fields for the film identifiers, image identifiers, image version, image counter and the sequence number are designated as primary keys for sorting the image reference table as well as for establishing relationships among the image reference table and other tables and/or databases.

The image table 2406d may include information such as: image identifiers, image version, original owner, image type, resolution, maximum width, maximum height, access control, purge flags, branding identifier, creation status, creation date, modification date, reference counter, external film identifier, external user name, external filename, source identifier, rotation angle, archived flag, copyright flag, and attribute flag. In one implementation, fields for the image identifiers and image version are designated as primary keys for sorting the image table as well as for establishing relationships among the image table and other tables and/or databases.

The film attributes table may include information such as: film identifiers, image version, attribute type, and attribute string. In one implementation, fields for the film identifiers, image version, and attribute type are designated as primary keys for sorting and searching the film attributes table as well as for establishing relationships among the film attributes table and other tables and/or databases.

The image attributes table may include information such as: image identifiers, image version, attribute type and attribute string. In one implementation, fields for the image identifiers, image version and attribute type are designated as primary keys for sorting the image attributes table as well as for establishing relationships among the image attributes table and other tables and/or databases.

The film identification counter may include information such as: hash value and identifier counter. In one implementation, the field for the hash value is designated as the primary key for sorting the film identification table as well as for establishing relationships among the film identification table and other tables and/or databases.

The image copyright table may include information such as: image identifiers, image version, line number, and copyright information. In one implementation, fields for the image

identifiers, image version and line number are designated as primary keys for sorting the image copyright table as well as for establishing relationships among the image copyright table and other tables and/or databases.

The account information table 2408a may include information such as: account name (e.g., screen name), a unique identifier (e.g., hash value) associated with the account name, image storage space used, image storage space available, account type, billing cycle, account creation date, modification date, account preferences, number of rolls, number of albums, number of shared accounts (e.g., buddies), notification flags concerning offensive content, pertinent dates (e.g., last login, last roll received, last image uploaded, last album shared, last mail received, last mail sent, welcome kit received), and/or other information. In one implementation the account name is designated as the primary key for sorting the account information table as well as for establishing relationships among the account information table and other tables and/or databases.

Referring to Fig. 5, the computer system 100 described above in connection with Figs. 1-4 operates according to a procedure 500. The procedure 500 may be implemented by any suitable type of hardware (e.g., device, computer, computer system, equipment, component); software (e.g., program, application, instructions, code); storage medium (e.g., disk, external memory, internal memory, propagated signal); or combination thereof. In general, the procedure 500 relates to storing digital imaged in a distributed fashion within the computer system 100.

Initially, a new image service account is created (step 505). In general, the host system 20 creates image service account for subscribers. In one implementation, the host system 20 includes an OSP host complex 230 that receives subscriptions from subscribers and offers images services as part of a package of online service available to subscribers. Typically, the host system 20 (e.g., account manager 244) automatically creates an image services account for a subscriber at the time of OSP subscription. In other implementations, the host system 20 may offer image services to customers of a film-processing partner regardless of the customers' ISP. In such implementations, the host system 20 (e.g., account manager 244) typically will prompt the subscriber to enter account information through a user interface.

Creating an image service account may involve automatically generating certain account information (i.e., fields). For example, typical account information may include but is

not limited to account name (e.g., screen name), a unique identifier (e.g., hash value) associated with the account name, image storage space used, image storage space available, account type, billing cycle, account creation date, modification date, account preferences, number of rolls, number of albums, number of shared accounts (e.g., buddies), notification flags concerning offensive content, pertinent dates (e.g., last login, last roll received, last image uploaded, last album shared, last mail received, last mail sent, welcome kit received), and/or other personal information (e.g., mailing address, phone number, instant message contacts, e-mail contacts). In general, the host system 20 (e.g., account manager 244) may be configured to automatically detect and update information to populate these fields with little or no input required by the subscriber.

Creating an image service account also may involve storing account information in a distributed fashion. In general, the account information is maintained on the host system 20 in massive storage facilities. In one implementation, the host system 20 includes an image farm 240 and an account manager 244 for populating the image farm 240. The image farm 240 includes a plurality of image farm databases 2404 for storing account information in account tables 2408. Each of the image farm databases 2404 may be assigned to one of a plurality of storage space groups (e.g., buckets). Account information is relatively uniformly distributed to a particular storage space group, i.e. bucket, and stored in an appropriate table (e.g., account information table 2408a) in one of the bucketed image farm databases 2404.

Distributing the account information to an appropriate storage space group may include determining a unique hash value by encoding the account information. For example, the account name (e.g., screen name) associated with a subscriber may be encoded using a proprietary OSP hashing code. In one implementation, the host system 20 (e.g., account manager 244) applies a hashing code that involves multiplying the individual ASCII values of the characters in the screen name by certain prime numbers, summing them, and returning a value modulo 64K (i.e., a 16-bit unsigned integer value between 0-65535). The returned hash values are mapped in round-robin fashion to the storage space groups, i.e., buckets, according to Table 1.

Hash Value	Bucket
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15
16	0
17	1
18	2
...	...
65535	15

**Table 1**

- 5 At any time after creating the image service account (step 505), the host system 20 (e.g., account manager 244) may update and/or delete the image service account. In general, the subscriber may actively request the host system 20 to update or cancel an image service account by logging in and requesting an update or cancellation. The host system 20 also may

delete a subscriber's account due to inactivity. For example, the host system 20 may remove an image service account that has been idle if a subscriber has not developed a roll within the last X (e.g., 6) months, uploaded an image within the last X months, shared an image (e.g., by e-mail or by sharing an album) within the last X months, or logged in the image service feature within the last X months. Of course, the timing on these activities should be flexible to allow business to adjust as market conditions demand. In addition, the host system 20 may delete a subscriber's account due to a policy violation. For example, sharing inappropriate digital images may result in cancellation of a subscriber's image service account, as discussed in greater detail below.

Following the creation of an image service account (step 505), one or more digital images associated with a subscriber are received (step 510). In general, a host system 20 receives digital images associated with a subscriber. In one implementation, the host system 20 includes a film handler 242 configured to receive digital images from a subscriber or from a third party (e.g., film developer) associated with the subscriber. For example, the film handler 242 may receive one or a collection of digital images from a subscriber or from a third party (e.g., film developer). For images received from a third party, each of the digital images may have a preset maximum resolution (e.g., 768x512 pixels) for incoming rolls. For subscriber-uploaded images, the size of the digital images may be variable with no preset maximum resolution.

After one or more digital images are received (step 510), a first storage facility for storing a digital image is identified (step 515). In general, the host system 20 identifies a first storage facility for storing a digital image. In one implementation, the host system 20 includes an image storage server system 2410 having a plurality of image storage servers 2412 configured to store digital images. The host system 20 further includes a film handler 242 that identifies one of the plurality of image storage servers 2412 as being available to store a digital image. For example, the film handler 242 may access a configuration file that contains a complete list of image storage servers 2412 and selects particular image storage servers 2414 in round-robin fashion. After selecting a particular image storage server 2412, the film handler 242 may communicate with the image write server 2414 and confirm that the image storage sever 2412 is capable of storing the digital image.

After the first storage facility has been identified (step 515), a directory within the first storage facility is identified for storing a digital image (step 520). In general, the host system 20 identifies a directory in the first storage facility. In one implementation, the host system 20 includes a plurality of image storage servers 2414, each image storage server 2414 having a directory structure for storing digital images. Typically, the directory structure for an image storage server 2412 will include several tiers identifying the storage facility, a directory, and several subdirectories. The host system 20 may include a film handler 242 that identifies one of the plurality of directories within a particular image storage server 2412 as being available to store a digital image. For example, the film handler 242 may access a configuration file that contains a complete list of first tier directories associated with each image storage server 2412. The film handler 242 may select particular first tier directories in the image storage server 2414 in round-robin fashion. After selecting a particular first tier directory, the film handler 242 may communicate with the image write server 2414 and confirm that the particular image storage sever 2412 is capable of storing the digital image within the particular first tier directory.

After a particular directory with the storage facility has been identified (step 520), a first image identifier is generated (step 525). In general, the host system 20 generates a first image identifier associated with the identified storage facility and directory. In one implementation, the host system 20 includes an image write server 2414 that generates a first image identifier corresponding to the identified storage facility and directory. One example, of a first image identifier is an 8 hexadecimal (32 bit) character string in which the first three hexadecimal characters correspond to the storage facility, the next three hexadecimal characters correspond to the directory, and the last two hexadecimal characters correspond to an encryption method. Typically, the groups of hexadecimal characters corresponding to the storage facility, directory, and encryption are coded. To illustrate, an example of a first image identifier (image\_id\_p1) is FEDCBA98. The characters FED correspond to a particular storage facility, for example, the image storage server 2412 named ygppics-d01.blue.isp.com. In this illustration, the characters CBA correspond to a directory within the particular storage facility, for example, the main storage directory 010. The characters 98 correspond to an encryption method, for example, the MD5 algorithm or DEC algorithm.

Next, a second image identifier is generated (step 530). In general, the host system 20 generates a second image identifier for determining particular subdirectories of the identified storage facility. In one implementation, the host system 20 includes an image write server 2414 that generates a second image identifier using a random number generator. One example of a second image identifier is a randomly-generated 8 hexadecimal (32 bit) character string. To illustrate, an example of a second image identifier (image\_id\_p2) is 76543210. Typically, the image write server 2414 will generate random numbers.

Then, a unique hash value is generated by encrypting the first and second image identifiers (step 535). In general, the host system 20 encrypts the first and second image identifiers and returns a hash value. In one implementation, the host system 20 includes an image write server 2414 that encrypts the first and second image identifiers according to a selected encryption method (e.g., MD5, DEC). One example of an encryption method involves applying the MD5 algorithm to the first and second image identifiers represented as a single character string (e.g., a single 16 hexadecimal character string). To illustrate according to the above example, the first image identifier (image\_id\_p1) and the second image identifier (image\_id\_p2) can be represented as the 16 hexadecimal character string, e.g., FEDCBA9876543210. A unique hash value is returned when the MD5 algorithm is applied to this character string.

In general, encrypting the first and second image identifiers returns an expanded hash value having more characters than the first and second image identifiers. In one implementation, the first and second image identifiers each include 8 hexadecimal characters and the returned hash value is a unique 32 hexadecimal character string value. To illustrate, assume that a trivial encryption method involves simply inserting random characters into a 16 hexadecimal characters string (image\_id\_p1 image\_id\_p2) to create a 32 hexadecimal character string. In such as case, the result of encrypting the first image identifier (image\_id\_p1) and the second image identifier (image\_id\_p2) could be the unique image identifier (ID): F1E2D3C4B5A69788796A5B4C3D2E1F0. Of course, in practical implementations, more elaborate hashing algorithms, such as the MD5 algorithm and/or the DES algorithm, may be utilized so that the resulting hash value gives no clear indication of the type of encryption algorithm being employed.

Next, a storage path is identified using the first and second image identifiers and the unique hash value (step 540). In general, the host system 20 identifies a storage path for the digital image using the first and second image identifiers and the unique has value. In one implementation, the host system 20 includes an image write server 2414 configured to extract storage path information from the from the first and second image identifiers. For example, the image writer server 2414 may extract storage path information corresponding to the tiers of a directory structure of an image storage server 2412.

As described above, each storage server 2412 may store digital images according to the following directory structure:

host:/data/AAA/BB/CC/DD/EE/ID\_N.type

In this pneumonic directory structure, host identifies a particular image storage server 2412 (root directory) for storing a digital image, data identifies a static directory of the image storage server 2412, AAA is a three hexadecimal character directory identifier denoting a main storage directory of the image storage server 2412, BB, CC, DD, and EE are each two hexadecimal character subdirectories in the image storage server 2412, ID is an encrypted 32 hexadecimal character image identifier that provides a file name and contains location information, N is the largest pixel dimension of an image (width or height), and type identifies image type (e.g., jpg).

Identifying the storage path may include extracting storage path information from the first and second image identifiers. In one example, the first image identifier (image\_id\_p1) is FEDCBA98 and the second image identifier is 76543210. These characters represent may represent an encoded path (e.g., href) of a storage location in the image storage server 2412.

The host system 20 (e.g., image write server 2414) may extract path information from the first and second image identifiers as follows: FED = host, CBA = AAA, 98 = encryption code (encryption used to obtain the hash value), 76 = BB, 54 = CC, 32 = DD, and 10 = EE.

Identifying the storage path may involve translating the extracted path information using a decoder (e.g., a look up table). For example, the decoder may translate the host characters FED into the sever name ygppics-d01.blue.isp.com, the AAA characters CBA into the main directory 101, the BB characters 76 into the first subdirectory 05, the CC characters

54 into the second directory 72, the DD characters 32 into the third subdirectory FA, and the EE characters 10 into the fourth subdirectory CB.

Identifying the storage path further may include using the unique hash value as a filename. As described above, encrypting the first and second image values returns a unique  
 5 filename. This unique hash value may be used as a unique image identifier (ID) that defines a unique file name and location for the image data. In one example, assuming the largest pixel value is 96 and the digital image is jpg type, the resulting path (href) in this case would be:

[http://ygppics-d01.blue.aol.com/data/010/05/72/FA/CB/hashvalue\\_96.jpg](http://ygppics-d01.blue.aol.com/data/010/05/72/FA/CB/hashvalue_96.jpg).

From the above, it is evident that the determination of a storage path is independent of typical image identifiers (e.g., account name, roll, date received). It follows, therefore, that related digital images (e.g., same account name, same roll, same date received) will have unrelated storage paths. For example, digital images from the same roll generally will have  
 15 storage paths corresponding to different image storage servers 2412, different directories, and different subdirectories. As such, the failure of one storage facility should have very little impact on the service of a particular subscriber. Furthermore, in the unlikely event that the storage location of a particular image is hacked, the decoded location of one image will give no indication of the storage location of any other related images. Consequently, the overall  
 20 security of the image service is improved.

After the storage path is identified (step 540), the digital image is stored at the identified storage path (step 545). In general, the host system 20 stores a digital image at the identified storage path. In one implementation, the host system 20 includes an image write server 2414 that receives a digital image from a film handler 242 and stores the digital image  
 25 in a particular image storage server 2412 at the identified storage path. Because the storage location is derived from a unique hash value, the storage location will be unique to that particular digital image.

Storing the digital image may include generating and storing lower resolution thumbnails. In general, the host system 20 generates the lower resolution thumbnails. In one  
 30 implementation, the host system 20 includes an image write server 2414 that uses a thumbnail function that creates thumbnails by ripping lower resolution images directly from a digital

image in the directory in which it was stored. The thumbnail function may rip and return a 96x64 pixel thumbnail or 160x107 pixel thumbnail, for example. The thumbnails will be named for the passed in ID and the resolution. Therefore, the lower resolution thumbnails may be stored in the exact same path as the path of the original.

5 After the digital image has been stored (step 545), a second storage facility is identified for storing metadata describing the digital image (step 550). In one implementation, the host system 20 includes a film handler 242 configured to identify an appropriate image farm database 2404 separate from the image storage server system 2410. Identifying an appropriate storage space may include determining a storage space group containing the appropriate image farm database 2404 from account information associated with the stored digital image. For example, the account name (e.g., screen name) associated with a subscriber may be encoded using a proprietary OSP hashing code. In one implementation, the film handler 242 applies the same proprietary hashing code to the subscriber's screen name as described above and is mapped by the distribution server 2402 to the appropriate storage space group, i.e., bucket, containing the image farm database 2404.

10 Then, metadata describing a digital image is stored in the second storage facility (step 555). In general, the host system 20 stores the metadata describing a digital image. In one implementation, the host system 20 includes bucketed image farm databases 2404 including film and image tables 2406 and account tables 2408 for storing metadata. The image farm databases 2404 and tables are separate from the image storage server system 2410, but generally contain metadata describing and pointing to the digital images stored in the image storage server system.

15 Storing metadata may include storing the first and second image identifiers in appropriate film and image tables 2406 (e.g., film reference table 2406a, film table 2406b, image reference table 2406c, image table 2406d). As described above, the stored first and second image identifiers may be used to identify a unique filename and storage location of a digital image stored in the image storage server system 2410. Storing metadata also may include automatically creating new records, generating certain metadata (e.g., unique film identifier, copyright attribute) and storing the metadata in appropriate film and image tables 2406. For example, the film handler 242 or a stored procedure in the film and image tables 2406 may be used to generate a unique film identifier. Storing metadata also may include

updating account information stored in account tables 2408. For example, when a new digital image is stored, the film handler 242 may update the subscriber's number of rolls counter or number of albums counter in the account table 2408a.

Finally, access to digital images and metadata describing the digital images is provided (step 560). In general, the host system 20 provides access to digital images and metadata to subscribers and/or other elements within the host system 20, such as the account manager 244 and the monitoring system 246, for example. In one implementation, the image farm databases 2404 are in communication with and capable of providing information (e.g., path information) to various elements of the OSP host complex 230 of Fig. 2 including the account manager 244.

Providing access to digital images and metadata describing the digital images may include allowing subscribers to view one or more digital images. In general, the film and image tables 2406 are searchable by various criteria (e.g., primary keys) including screen name. For example, by passing a subscriber's screen name to the film and image tables 2406, the account manager 244 can retrieve image records and/or film records associated with the subscriber's account. Typically, the retrieved records will include image identifiers referencing one or more stored digital images. The account manager 244 can use the image identifiers (e.g., image\_id\_p1 and image\_id\_p2) to determine the storage path of a particular digital image stored in a particular image storage server 2412 within the image storage server system 2410, as described above. By navigating to the appropriate storage path in the image storage sever system 2410, the account manager 244 can retrieve and display a particular digital image to a subscriber.

Providing access to digital images and metadata also may include allowing subscribers to list and view all rolls for a screen name passed as a search parameter, to list and view all albums for a screen name passed as a search parameter, to list and view all buddy albums for the screen name passed as a parameter, and to list all images for a film identifier passed as a parameter. After listing and viewing certain digital images, rolls, and albums a subscriber may edit and/or delete the digital images, rolls, and albums.

Providing access to digital images and metadata also may include allowing subscribers to share a roll or album. In general, the host system 20 allows subscribers to share digital images. In one implementation, the account manager 244 shares a digital image by taking

film identifier and buddy screen name as an argument, locating a film reference table 2406a in the owner's storage space group (e.g., bucket), creating a new record in the owner's film reference table, accessing a film reference table 2406a in the buddy's storage space group (e.g., bucket), creating a record in the buddy's film reference table, and incrementing the number of buddies field in the account information table 2408a in the buddy's storage space group.

Referring to Fig. 6, the computer system 100 described above in connection with Figs. 1-4 operates according to a procedure 600. The procedure 600 may be implemented by any suitable type of hardware (e.g., device, computer, computer system, equipment, component); software (e.g., program, application, instructions, code); storage medium (e.g., disk, external memory, internal memory, propagated signal); or combination thereof. In general, the procedure 600 relates to monitoring stored digital images within the computer system 100.

As described above, subscribers are able to share digital images with each other. Typically, an owner will create an online album of digital images and then invite one or more other subscribers to view the online album. An invitation may be extended to the one or more other subscribers by an e-mail message. For example, the owner may assemble a list of invited subscribers and click an "invite" button. Elements of the host system 20 (e.g., mail system 238, account manager 244) may react to the owner's invitation by sending an e-mail message to each of the invited subscribers. The e-mail message may include a password and storage path to the shared digital images.

While this important feature of the image service allows families and friends to share life events captured with digital images, there exists potential for abuse. Namely, a subscriber may be invited to view a digital image that he or she finds offensive. Typically, this situation arises when an owner obtains a list of e-mail addresses through dubious means and mass invites the group to view adult-oriented material.

To protect subscribers, the host system 20 is configured to receive notifications from subscribers regarding offensive content. In one implementation, the host system 20 includes a monitoring system 246 configured to receive notifications of offending content from subscribers and take appropriate actions in response.

Initially, a notification regarding offensive content is received (step 605). In general the host system 20 provides a user interface allowing subscribers transmit notifications. For example, shared albums may include a "notify" button that allows an invited subscriber to

object to one or more digital images. In one implementation, when a subscriber clicks the notify button, an e-mail is sent to the monitoring system 246.

Next, an investigation report is generated (step 610). In general, the host system 20 generates the investigation report. In one implementation, the host system 20 includes a monitoring system 246 configured to compile investigation reports in response to notifications from subscribers. Generating an investigation report may involve compiling information regarding the offending digital image. Some information may be supplied by the complaining subscriber and some information may be automatically generated by the monitoring system 246. Report information may include: the photo ID and the album ID of the picture album, the member created picture album name and the name, if any, of the individual photos, the screen name of the picture album owner/creator, an indicator if the album owner is a member of the OSP or another type of member, the country code of the picture album owner, the date and time the picture album was created/edited, the file size of the reported photo ID for the low resolution image size, the date and time that the report was sent to the image monitoring system 246, the screen name of the member reporting the picture album, client information (e.g., platform, client version) of the member reporting the album, the violation category of the offending image, comments from the complaining subscriber, and the list of shared members for the reported picture album.

Within a very short period of time (e.g., 5 minutes) after the notification is received generated, a storage location associated with the complaining subscriber is identified (step 615). In general, the host system 20 an appropriate database (e.g., image farm database 2404) containing metadata associated with the complaining subscriber. Identifying the appropriated database may include determining a storage space group containing the appropriate image farm database 2404 from account information associated with the complaining subscriber. For example, the account name (e.g., screen name) associated with the complaining subscriber may be encoded using a proprietary OSP hashing code. In one implementation, the monitoring system 246 applies the same proprietary hashing code to the subscriber's screen name as described above and is mapped by the distribution server 2402 to the appropriate storage space group, i.e., bucket, containing the image farm database 2404.

After the storage location is identified (step 615), metadata associated with the complaining subscriber is modified to block the complaining subscriber from reviewing the

offending image (step 620). In general, the host system 20 modifies metadata associated with the complaining subscriber and metadata associated with the offending digital image. In one implementation, the monitoring system 246 finds a film reference table 2406a corresponding to the offending digital image in complaining member's bucket. The monitoring system 246 invokes an unsharing procedure that takes the complaining subscriber's screen name and film identifier corresponding to the offending digital image as parameters and searches film and image tables 2406 and account tables 2408 in the complaining subscriber's bucket. The monitoring system 246 removes metadata corresponding to the offending digital image from tables (e.g., account information table 2406a) associated with the complaining subscriber and removes metadata corresponding to the complaining subscriber from tables (e.g., film reference table 2406a) associated with the offending digital image.

Then, a storage location associated with the owner of the offending digital image is identified (step 625). In general, the host system 20 identifies an appropriate database (e.g., image farm database 2404) containing metadata associated with the owner of the offending digital image. Identifying the appropriated database may include determining an account name (e.g., screen name) associated with the owner of the offending digital image. The owner's screen name may be obtained from the film reference table 2406a in the complaining subscriber's bucket, for example. Identifying the appropriate database also may include determining a storage space group containing the appropriate image farm database 2404 from account information associated with the complaining subscriber. For example, the account name (e.g., screen name) associated with the offending digital image may be encoded using a proprietary OSP hashing code. In one implementation, the monitoring system 246 applies the same proprietary hashing code to the owner's screen name as described above and is mapped by the distribution server 2402 to the appropriate storage space group, i.e., bucket, containing the image farm database 2404.

After the storage location associated with the owner of the offending digital image is identified (step 625), metadata associated with the owner is modified to block the owner from sharing the digital image with the subscriber (step 630). In general, the host system 20 modifies metadata associated with the owner and metadata associated with the offending digital image. In one implementation, the monitoring system 246 finds a film reference table 2406a corresponding to the offending digital image in the owner's bucket. The monitoring

system 246 invokes an unsharing procedure that takes the owner's screen name and film identifier corresponding to the offending digital image as parameters and searches film and image tables 2406 and account tables 2408 in the owner's bucket. The monitoring system 246 then removes metadata corresponding to the complaining subscriber from tables (e.g., film reference table 2406a) associated with the offending digital image and inserts metadata (e.g., purge flag) into tables (e.g., film reference table 2406a, account information table 2408a) indicating that the offending digital image is under investigation. At this stage of the procedure 600, the complaining subscriber can no longer access the offending digital image. However, reporting a digital image or album will have no effect on the ability of other subscribers on the roster of a shared album to view the digital image or album until after an administrative review is completed.

Next, a separate storage location containing the digital image is identified (step 635). In general, the host system 20 identifies the storage location of digital images. In general, the film and image tables 2406 are searchable by various criteria (e.g., primary keys) including screen name. For example, by passing a complaining subscriber's and/or owner's screen name to the film and image tables 2406, the monitoring system 246 can retrieve image records and/or film records associated with the offending digital image. Typically, the retrieved records will include image identifiers referencing one or more stored digital images. The monitoring system 246 can use the image identifiers (e.g., image\_id\_p1 and image\_id\_p2) to determine the storage path of a particular digital image stored in a particular image storage server 2412 within the image storage server system 2410.

Then, the offending digital image is retrieved (step 640). In general, the host system 20 retrieves the offending digital images. In one implementation, the host system 20 includes a monitoring system 246 configured to navigate to the appropriate storage path in the image storage sever system 2410.

Next, the offending digital image is reviewed for compliance with the terms of the image service (step 645). In general, the host system 20 displays the offending digital image for review. In one implementation, the monitoring system 246 displays a particular offending digital image to an OSP administrator for review. Typically, the terms of service (e.g., policies) are established by administrators of the OSP host complex 230 and regulate the type of content that is acceptable for sharing among subscribers.

Reviewing a digital image may involve analyzing an investigation report associated with the offending digital image. In general, an investigation report associated with an offending digital image is stored in a database queue to await action by an OSP administrator. Typically, several OSP administrators will be available to conduct reviews. Investigation reports are routed to an appropriate OSP administrator based on the country code of the owner of the digital image.

Reviewing a digital image may involve classifying an investigation report. In general, the investigation reports may be categorized according to the reported type of violation and review status. For example, all new reports may be classified as a "general" violation. As reports are reviewed, depending on the content of the page, the report may be designated as "acceptable" or "unacceptable." In addition, a report may be designated as non-garden variety "NGV" violation requiring further review. Further, the OSP administrator reviewing NGV reports may mark reports as pending legal review "PLR" in extreme cases. In general, the OSP administrators may include different specialists (e.g., legal personnel) for reviewing different categories of violations.

Finally, appropriate action is taken based on the review of the digital image (step 650). In general, the host system 20 takes appropriate action based on the review of the digital image by the OSP administrator. In one implementation, the host system 20 includes a monitoring system 246 configured to take different levels of action depending upon the nature of the offending content. Such actions may range, for example, from blocking the complaining subscriber from reviewing the offending digital image to deleting the digital image, canceling the offender's account, and alerting law enforcement authorities.

If after reviewing the digital image, the OSP administrator determines that there is no violation, the OSP administrator will indicate that the digital image is "acceptable." While the complaining subscriber will have no further access to the particular image, other subscribers sharing the digital image will be unaffected.

If the digital image is in violation of the terms of service, the OSP administrator indicates that the digital image is "unacceptable." In this case, the monitoring system 246 will send a warning e-mail to the owner of the unacceptable digital image, and document the violation in the owners account history. For milder violations, the monitoring system 246 may unshare the digital image. This function will allow the owner to still have access to their

digital images, but no longer allow any other subscribers to view the album. Depending on the severity of the violation, the monitoring system 246 also may delete the digital image and metadata describing the digital image from all table and/or terminate the owner's account.

5 In cases where the content of the digital image warrants further review, the OSP administrator may forward the digital image and its associated report to a NGV specialist. The NGV specialist may make a final determination as to whether the digital image is acceptable or unacceptable.

10 In extreme cases, where the content of the digital image may be illegal, the OSP administrator may forward the digital image and its associated report to legal personnel. The monitoring system 246 may block the violator's entire picture album from being viewed by anyone except OSP administrators and legal personnel. The legal review specialist may make a final determination as to whether the digital image is acceptable or unacceptable. In some cases, the legal specialist may inform law enforcement.

15 A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made and that other implementations are within the scope of the following claims. For example, if a more elaborate method of preventing unauthorized users from walking the href to find image paths is needed, an encryption algorithm can be applied to the first and second image identifiers. The tradeoff is the time it takes to decrypt the image identifiers on every image view.

10007596-12001